

Benefits of Using Explicit Ground-Plane Information for Grid-based Urban Environment Modeling

Jens Rieken, Richard Matthaei, and Markus Maurer
Technische Universität Braunschweig
Institute of Control Engineering
38106 Braunschweig, Germany
Email: <rieken,matthaei,maurer>@ifr.ing.tu-bs.de

Abstract—Future advanced driver assistant systems and automated driving put high demands on the environmental perception especially in urban environments. Major tasks are the perception of static and dynamic elements along with the drivable area and road structures. Although these tasks can be done without an explicit representation of the ground surface, evaluations with real-world sensor data have shown that many failures in sensor data interpretation result from inappropriate ground point filtering, e.g. false negative or false positive detection of objects. That means, the ground plane estimation becomes a key-task in urban environment modeling assuming a reliable perception of the local vehicle’s surroundings is required. In this paper we thus present the enhancements of the environment modeling that can be derived from an explicit representation of the ground surface. We propose the combination of channel-wise ground classification with a grid-based representation of the ground height to cope with spatial false measurements and exploit inter-channel dependency. Partial occlusions by other road users are handled in an efficient way. Integration of the ground data is shown by a pitch angle estimation and a curb detection module. The algorithms run in real-time on a standard PC and are evaluated with real sensor data.

I. INTRODUCTION

The development of advanced driver assistance systems (ADAS) and automated road vehicles has high demands on the environmental perception of the vehicle. Especially for inner-city scenarios, which are addressed e.g. by the project *Stadtpilot* [1], a complete representation of the vehicles’ surroundings is necessary to provide a safe and comfortable driving experience.

That means, the host vehicle has to perceive features of the ground surface as well as elevated elements in order to determine the course of the lanes and detect obstacles that would potentially collide with the host vehicle. Among these are dynamic objects, e.g. vehicles and vulnerable road users (VRUs), and static obstacles on and at the side of the road. On highway-like roads, the course of the lanes can be determined from ground features, e.g. lane markings or textural changes of the ground surface. But when advancing into inner-city streets, also curbs and elevated objects provide relevant information about the boundary of the roads and drivable areas.

These perception tasks mentioned above benefit from a stable and reliable detection and explicit representation of

the road surface. Evaluation with real-world sensor data have shown that systematic errors arise if the ground height around the host vehicle and its environment is not regarded. This results in either *false positive* objects (ghost objects from ground measurements) or *false negative* objects (overseen objects due to false classification as ground measurements).

When knowing the ground surface around the host vehicle, lane marking hypotheses can be validated based on the ground plane information. The height of obstacles is measured against the ground height below that obstacle, thus correct object height estimation as well as the detection of underdrivability are coupled to the ground plane. The knowledge of the ground plane enables us to extract curb features, which are characterized by a sudden moderate height change, in a simple way. The curb data provide additional features to road fusion algorithms. Additionally, the pitch and roll angle of the host vehicle in relation to the ground plane are relevant for image processing algorithms and classification tasks of other sensors. These angles can be determined from a ground plane estimation as well.

Although suggested by its name and often assumed when filtering for ground measurements (e.g. [2], summary in [3]), the ground plane is usually not a plane in terms of a mathematical definition. Locally plane-like, the ground area might contain slope changes over larger distances. As these slope changes have similar influence as pitch and roll angle changes, the ground surface model has to be able to deal with these slope changes in an appropriate way.

In this paper, we propose a grid-based algorithm for ground plane estimation which provides locally consistent results and which is capable of dealing with local slope changes up to a certain extent, defined by a maximum allowed slope. We demonstrate the benefit of a grid-based ground plane representation compared to a more simple channel-based classification of ground measurements. The practical usage of an available ground plane is shown by example applications: A local pitch angle estimation and the robust extraction of curb features along with an improved height classification of sensor measurements. The additional curb features are used to extend our road fusion module, already published in [4], to further

stabilize the road extraction process.

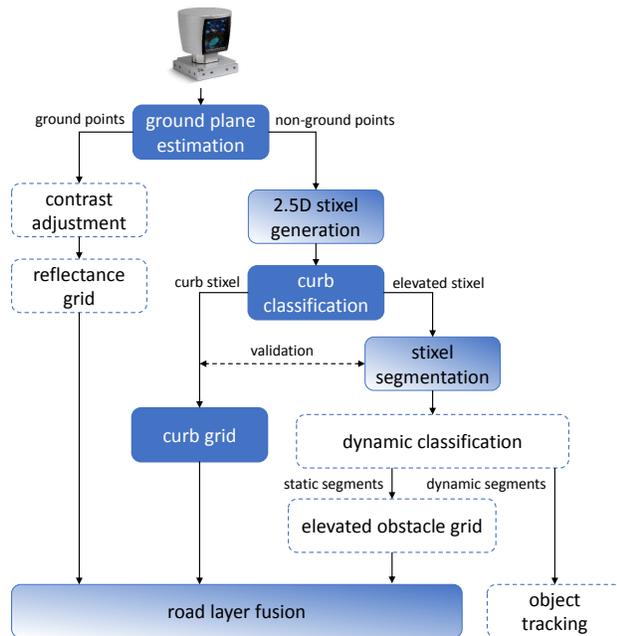


Fig. 1: Overview of the overall sensor data processing chain for the perception tasks. Colored modules are addressed by this contribution, partly-colored modules are extended in order to handle new features, grayed modules are out of scope.

This contribution is organized based on the signal data flow shown in Figure 1. Modules addressed by this paper are highlighted in blue. Other modules are mandatory for the overall perception task of the automated vehicle, but are outside the scope of this paper and are addressed by other publications of our research group. After the review of related work in Section II, the ground plane estimation is introduced in Section III. Based on these results, we present our pitch angle calculation based on the estimated ground plane in section IV. The adoptions of the stixel generation and segmentation step as well as the curb classification module are presented in Section V. These features are finally used to extend our road fusion framework (Section VI). We close with a conclusion and outlook to further work in Section VII.

II. RELATED WORK

In this paper, we refer to the term *ground* as that area which is reachable by the host vehicle without passing discontinuities, defined by a maximum allowed slope of the ground surface. In the following section, a brief overview over established algorithms is given and the contribution of our approach and its usage for environment modeling is shown.

A. Literature review

The estimation of the ground surface around the host vehicle is a major topic in the context of automated driving as it describes the drivable area. Several sensor technologies are

used to retrieve this information, e.g. computer vision systems (mono and stereo cameras) and laser distance sensors. An extensive comparison of different approaches is given by the authors in [3] and [5]. Here, we will focus on the extraction from a high density 3D point cloud generated from the well known Velodyne LIDAR. The classification of ground points from those data was addressed during the DARPA Urban Challenge in 2007, as many teams used this type of laser scanner (e.g. [6] [7] [8] [9]) or accumulated a 3D point cloud from multiple single-layer laser scanners (e.g. [10]).

In general, existing algorithms can be subdivided into two groups. Algorithms of the first category generate a ground classification based on the relative position of measurements inside a defined vertical slice of the sensor data (commonly known as *channel*). Algorithms of the second group estimate ground information in the scan data of a group of channels up to the complete sensor scan. Examples for algorithms of both groups are given in the following paragraphs.

a) *Channel-based algorithms*: Channel-based classification algorithms compare the given measurements within a scan channel to identify those measurements which describe the ground surface. Several different approaches can be found in literature. Ground points can be detected by comparing the vertical displacement between the sensors' measurement inside a bounded area. By this, flat structures can be detected and separated from elevated targets (e.g. [8] [10] [11]). Slope-based algorithms compare the height data from subsequent measurements inside a channel in relation to the radial distance to detect slope changes and thus measurements from vertical structures. Beside direct angular calculation, adjacent beam comparison is applied to be more robust against high-frequency dynamic movements of the vehicle [6]. A similar algorithm was used in the first place by the Stadtpilot project as well (see [12]).

In contrast to the aforementioned approaches, the authors in [13] applied a line regression algorithm after clustering the sensor raw data within a 2.5D grid representation. The extracted line was used as the ground slope inside the channel. The measurements were classified afterwards regarding their distance to this 'ground line'.

Typically, these algorithms are used to separate ground points from measurements of elevated targets in order to remove the ground points from the sensor scan. This allows to detect elevated point clusters and drivable areas. An explicit modeling of the ground surface was not used in these approaches. The channel-based classification is sensitive to ground slope changes, but does not take the neighbor channels into account, which would lead to a more stable and consistent classification.

b) *Group-Based and Scan-based algorithms*: Other approaches use the explicit modeling of a plane function in order to classify points as part of the ground plane. The authors in [14] discretized the sensor data into distinct cells and applied a RANSAC fitting to the lowest z coordinate inside those cells, assuming these measurements to be part of the ground surface. The authors in [15] use a RANSAC algorithm as well,

but apply it to the nearly complete sensor point cloud. Only points above a certain height in the sensor reference frame were removed as they cannot be generated from ground measurements. Other approaches, e.g. [2], combine the channel-based classification with a plane fitting algorithm. A slope-based approach is used to determine possible ground points. In the second step, a plane is fitted into this point cloud in a least-squares manner.

As stated previously, these approaches assume the existence of a single ground plane around the vehicle (*flat world assumption*). If applied to non-planar areas, these algorithms lead to false approximations of the ground height and thus false measurement classification. False classification leads to failures in the object detection, e.g. ghost objects then classifying ground points as elevated targets or missing objects when classifying elevated objects as part of the ground plane.

B. Contribution of this paper

Many of the above mentioned algorithms focus on the classification of sensor data in order to remove ground points from further processing steps or to extract the drivable area from the vehicle's point of view. Channel-based and scan-based algorithms differ in their ability to adapt to slope changes or to generate a consistent ground surface estimation. Similar to the approach in [2] and [9], we propose the usage of a channel-based pre-classification combined with an area-based consistency constraint. By using a grid-based representation, we are not limited to using one plane function for the overall ground surface, but can adapt to slope changes up to a predefined extent. Furthermore, the grid-based representation allows for the compensation of partial occlusions from other traffic participants. The results of the ground surface estimation are used to correct the height information of object hypotheses and to classify the extent of vertical structures. This allows for a simple curb classification and enhances the results of our grid-based road fusion. Furthermore, the correction of the height component results in a higher effective range of view since less false classifications would appear at larger distances.

III. GROUND PLANE ESTIMATION

Our ground plane estimation consists of four major processing steps to be explained in the following paragraphs.

A. Algorithm description

a) *Profile-based classification according to [16]*: After sensor data acquisition and preprocessing to compensate sensor decalibration, the sensors' data is organized in a channel-like data structure. All measurements belonging to the same azimuth angle of the sensors' reference frame are sorted by their elevation angle in ascending order.

The first processing step performs a profile-based classification inside such a channel. Given a pre-defined threshold for the maximum allowed ground slope (change of height over distance), the slope angle between two subsequent measurement points inside a channel is calculated by comparing the cylindrical radial distance change against the detected height

value change. Measurements with a slope below the applied threshold are marked as ground point candidates, whereas other points are classified as measurements from elevated targets. The ground point candidates are validated by a profile criterion. If a previous measurement in the current channel was classified as an elevated target measurement, the slope must have a downward direction larger than the given threshold. By applying this rule, we prohibit elevated plateaus from being classified as ground points. The results of this step are shown in Figure 2, left. The major number of ground measurements is classified correctly. However, sensor noise and small irregularities on the pavement, e.g. grass between cobblestones, lead to spatial false classifications in some channels. At this processing step, the algorithm lacks an explicit model of the dependency between the neighbored channels.

b) *Cell-based aggregation*: Spatial false classifications can be filtered by exploiting inter-channel dependencies. Typically, the ground surface height does not change significantly between the channels, so consistent height hypothesis can be generated when regarding neighbor channels. This dependency is modeled via a grid-based representation of the ground height around the host vehicle. All ground-classified measurements are taken into account in this processing step and are assigned to a certain grid cell depending on their position. This leads to a clustering of ground points to specific grid cells. In a second step, the mean height of all points inside such a cell is then calculated and provided as a mean height of the ground plane at that specific area. The usage of a grid structure combines the advantages of both the scan-based and the channel-based algorithms. It provides a locally consistent ground height and adapts well to ground height changes inside the sensor range as it uses the profile-based classification of the first processing step as an input filter. The results from this processing step are shown in figure 3, left.

c) *Median filtering and gap filling*: A modified median filter is applied to this grid structure to fill spatial missing ground plane information and to further flatten the detected ground height. The filter algorithm calculates the median value of all available ground cells inside its filter kernel. This allows the median filter to be able to operate even on cells with partial missing ground plane information in its neighborhood. Furthermore, this algorithm is able to fill up missing ground cells, if a sufficient number of ground cells is found in its surrounding. The results of this processing step are shown in Figure 3. We chose a grid cell size of 1.0 m per dimension and a median filter kernel size of 9 cells. With these parameters, we can compensate missing measurements underneath vehicles on the road, but are robust against single-point failures from the profile-based classification.

d) *Final point classification*: With knowledge about the ground plane around the host vehicle, all sensor measurements of the current scan can be classified according to their specific distance to this plane. While points with *ground* classification have a reasonable small distance to this plane (currently a threshold of 0.1 m), we also classified points with a distance up to 0.25 m as *curb*. Points above this threshold are classified

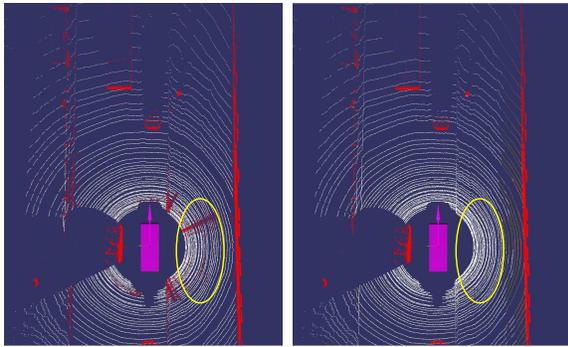


Fig. 2: Results of the profile-based classification (left) and the grid-based classification (right): ground points (white), elevated points (red, orange). The profile-based classifier yields spatial false classifications due to small irregularities and false measurements (yellow marking). The grid-based classifier provides a more robust classification.

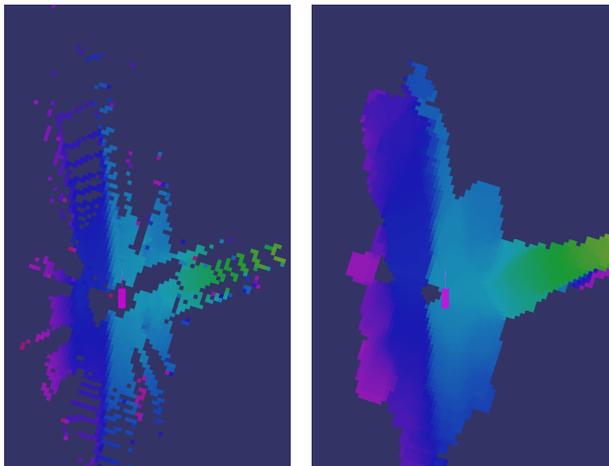


Fig. 3: Resulting ground height from grid-based approach before (left) and after (right) applying a median filter. The modified median filter fills cells with missing ground information and removes spatial outliers. The height is colored from low (green) over zero (cyan) to high (red).

as regular elevated targets.

This threshold-based decision yields good results so far, it leads to false classifications on vertical structures, e.g. house walls. As the ground plane cuts with the vertical structure, measurements with small distance to the ground plane on the vertical structure would be classified as curb as well. To overcome this issue, the number of non-ground measurements inside each grid cell is calculated. If the number of non-ground measurements relative to the number of ground measurements is above a pre-defined threshold, we assume a vertical structure inside that cell. Measurements with a distance-based curb classification inside such a cell will be classified as *uncertain curb*. This allows subsequent modules a proper handling of

such measurements based on their demands, e.g. the contrast adjustment or the reflectance grid (out of scope in this paper; for further information about these processing steps, see [16]).

B. Computation performance

The applied algorithms are computationally feasible. Typical execution times, measured on an Intel i7-4770S Quadcore processor, are in a range of 10 – 14 ms. The execution time of the grid-based filtering scales with the total number of ground points classified in the first processing step and depends on the used grid cell size and cell count. The median filter is currently implemented in a straight-forward manner with $O(n^2)$ complexity. Thus, the choice of the median kernel size has significant influence on the overall performance. By using algorithmic optimizations (e.g. [17]), the processing time of this filter step can be reduced.

IV. PITCH ANGLE EXTRACTION

With information about the ground plane around the vehicle, we can extract the relative angle between the sensors' mounting position and the ground plane. The angles in x and y direction of the vehicle reference frame (according to DIN 70000) are known as pitch and roll angle. These angles have huge influence on the mounting position of the sensor relative to the road surface and thus on the interpretation of sensor measurements, especially for measurements at larger distances. The extracted angles can be used to compensate high-frequent vehicle movements due to braking and acceleration maneuvers or road unevenness. Examples for the usage of the relative pitch and roll angles are given in [16], [18] and [19]. Especially on roads with a slope, the ground-relative angle differs from the absolute pitch and roll angles measured by accelerometers of an inertial measurement unit (see Subsection IV-B).

The extraction is currently limited to the pitch angle only, because in dense traffic scenes or when driving through house canyons, the detection range along the lateral vehicle axis is fairly limited. The roll angle is less important for the sensor data correction as the pitch angle is the dominant component. Thus, the next paragraphs will focus on the pitch angle extraction based on 3D Lidar data. The extraction algorithm in general can be applied to the roll angle in the same manner.

A. Extraction algorithm

The pitch angle is extracted from virtual scan lines following the longitudinal axis of the vehicle reference frame. The extraction range is limited to 25 m in both directions, as valid ground layer cells are most likely available inside this range. A line fitting algorithm is applied to the extracted cell values (ground heights) to get a linear function of ground plane with regard to a perpendicular-least-square manner [20]. The total pitch angle is calculated as a weighted mean from both lines with the total extracted number of cells as weighting factor. This approach assumes a constant slope profile of the underlying ground plane along the extraction line; a sudden slope change between the front and rear line would result in

an erroneous angle. This influence can be reduced by either adjusting the extraction range or, in addition to that, using another fitting function primitive, e.g. a parabolic function.

B. Results

The results of the pitch angle estimation based on the ground plane are shown in Figure 4. The extracted *ground-relative* pitch angle is plotted against the *absolute* pitch angle measured from an iTrace-F200 inertial measurement unit, which is used as a reference. The high-frequent component of the extracted angle which results from road surface irregularities and braking/acceleration of the host vehicle, follows the pitch angle from the IMU without a noticeable delay. A low-frequent offset between both plots becomes evident towards the end of the sequence. This offset, the difference between the absolute and the ground-relative pitch angle, results from a slope of the ground plane. The absolute ground height (green line) measured via GPS localization of the iTrace shows a downward slope of the road. The angular difference correlates with the changing ground height, which shows the correct estimation the relative pitch angle.

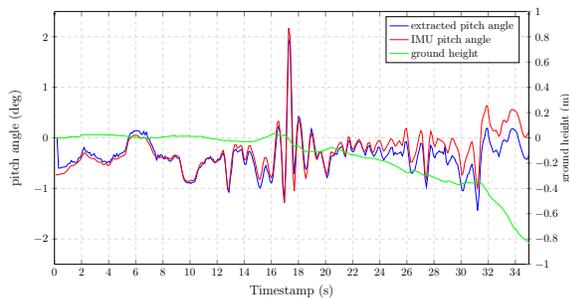


Fig. 4: Comparison of the extracted pitch angle with the values from iTrace system. The high-frequent component follows the values from the accelerometer. The low-frequent offset results from a change in the overall road surface slope. The ground height is plotted for comparison.

Figure 5 shows another measurement sequence taken on a highway segment around Braunschweig. The road segment is characterized by a slight slope over a large distance, resulting in a locally plain road surface. Without any disturbance due to acceleration/braking maneuvers and irregularities of the road surface, one expects a pitch angle around zero degrees. In deed, this is achieved by our extraction of the relative pitch angle: It has the same high-frequent peaks as measured from the IMU, but has a zero mean. The large-scale slope of the road is visible as a low-frequent offset between the IMU angle and the extraction angle; the delta is plotted as *extracted ground slope*. It matches with the deviation of the GPS/IMU height signal measured by the iTrace system.

V. PROCESSING OF STIXEL DATA

In a subsequent step of the ground plane estimation the non-ground measurements are processed. In order to reduce

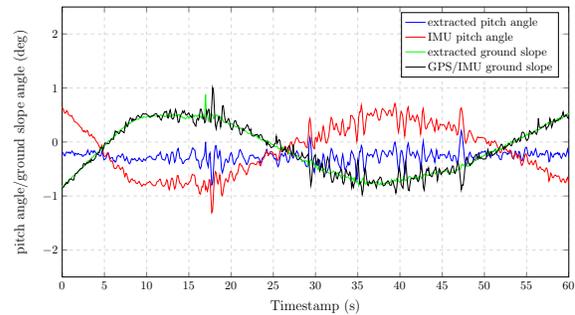


Fig. 5: Extracted ground slope compared to iTrace system estimated ground slope on a road segment with slope changes over time. The proposed algorithm extracts the expected near-zero pitch angle, while the IMU is influenced by the road slope. The resulting delta (extracted ground slope) is correlated with the deviation of the GPS height signal.

the total amount of measurements, a stixel representation is generated from the 3D point cloud data by projecting the data into a polar grid, as already published in [12]. In addition, we use the ground plane information at the respective cells to calculate the stixel height in relation to the ground height.

Extensive work has been done in the field of curb detection using different approaches, e.g. from computer vision [21] or 3D laser distance data [22]. A comprehensive summary and comparison of algorithms used on this topic is published in [5]. The approach presented in this paper uses a simple distance-to-ground calculation to classify points as curbs. Based on the ground-relative stixel height, we classify stixels with a height below a predefined threshold (currently 0.25 m) as *curb* and stixels above as regular targets. This closes a gap of our previous published approach in [16].

The segmentation used in this project is based on the well known algorithm presented by Dietmayer et al. [23] and exploits the channel-wise representation of the sensor data. This algorithm has been extended to handle curb-classified stixels differently during the segmentation step. Those stixels are neglected during the segmentation process, but are assigned to the according segment if they are inside the given geometric proximity threshold. By this, the segments will include all curb stixels fulfilling the segmentation criterion, but curb stixels are not able to connect segments, e.g. parking vehicles in front of a curb-like structure. The results of the curb classification and the segmentation step are shown in Figure 6.

Due to the sensors' resolution, detailed features like curbs can only be detected reliably inside a short area around the host vehicle, while elevated targets will be detectable at a large distance up to the sensors' viewing range. The different detection ranges lead to a late occupancy information for curb-like structures if they are accumulated inside the same Bayesian grid as elevated targets. As typical sensors 'oversee' curbs at a larger distance due to the growing gap between two vertical scan lines at larger distances, the Bayesian grid cells will be updated as 'free' up to the first elevated target.

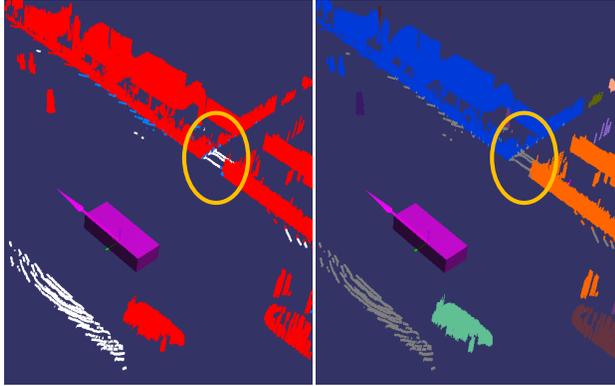


Fig. 6: Left: Results of the curb classification. Stixels are classified as *curb* (white), *uncertain curb* (curbs near elevated targets, blue) and *elevated* (red).

Right: Adapted segmentation process: Presence of curb stixel aborts the segmentation, which results in different segments at the elevated targets (yellow marking).

When finally detecting curb points, the relevant grid cell already contains a free state and will take some update steps to converge towards occupied state. This results in an unstable occupancy information on curbs and pavement areas around the host vehicle.

To avoid this issue, curbs and elevated targets are accumulated in two separated Bayesian grids based on their previous classification. The results are shown in Figure 7. The combined Bayesian layer shows unstable occupancy states in the area of the curb, while the curb layer contains stable information about the curb features and the distinct elevated target layer contains stable information about targets above the curb threshold. Both layers can be used as different features by e.g. the road area detection (see section VI).

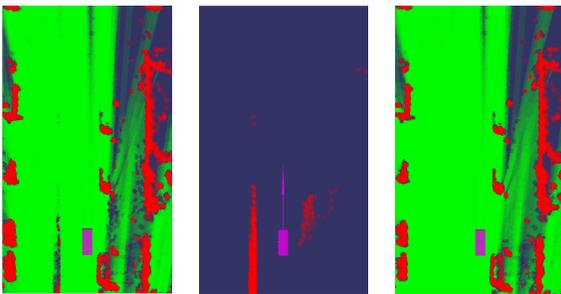


Fig. 7: Comparison of a default Bayesian grid (left) against separate layers for curbs and elevated targets (middle/right): Curb and elevated target features create high occupancy values in their respective layer while the combined layer contains unstable occupancy in the area of curb measurements (red: occupied, green: free; magenta box shows host vehicle position).

VI. INTEGRATION TO ROAD EXTRACTION

In [4] and [24] we presented an algorithm for road extraction. Our algorithm addresses the challenge of detecting road and lane geometries especially in urban environments which do not follow a certain model assumption, e.g. a clothoid model. Due to the big variety of environmental features in urban environments (e.g. elevated targets, lane markings, curbs) which all together define the road and lane boundaries, we presented in [16] and [24] an approach which fuses elevated targets and lane markings into one single grid-based representation. This step improves the robustness of the road detection algorithm referring to different local environments (e.g. side streets without lane markings or radial highways with lane markings but without building lines). However, we did not explicitly consider curb-like features so far.

The integration of the curb layer into our existing framework is shown in Figure 8. The curb information is fused with the information of elevated targets and lane markings to extend the low-level representation of those features which define the road course.

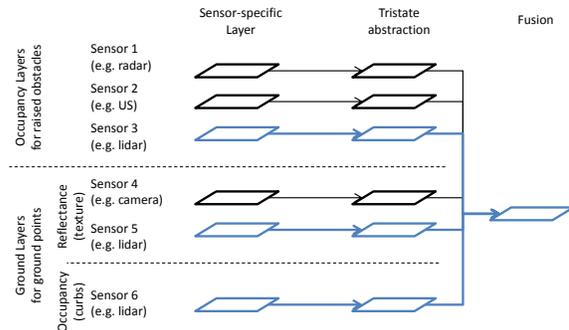


Fig. 8: The curb layer in the context of our multi-layer concept: Different features are stored in specialized layers, abstracted to a common representation and then fused per-cell by a rule-based framework. The resulting representation is used by e.g. our grid-based road extraction algorithm (blue parts are currently used, black ones can be additionally integrated).

In this framework, several features are stored in dedicated grid layers. The cell data type and the data processing algorithms can thus be optimized for the respective features.

The grid fusion algorithm combines the information stored in those different layer types. To be able to perform this fusion step, each cell of the feature layers is converted into a tristate value (free, unknown, occupied) by a feature-dependent algorithm. This values correspond to the presence (*occupied*) or absence (*free*) of a specific feature; the *unknown* state is used to label cells with inconclusive information. The different layer abstractions are then fused on a per-cell basis into a combined representation using a rule-based framework. The fused states of the cells provide information about the availability of features at the cells' positions. This includes

both information about the type of detected features and the representation of free and unknown areas.

The road extraction algorithm uses the type of occupancy stored in this fused representation to select suitable features for the determination of road and lane courses.

The fusion of curbs and the structure of the ground plane (reflectance layer) provides redundant information about the road borders because in many cases the material of the slightly elevated curbs differs from the material of the road surface (see Figure 9 and Figure 10). Nevertheless, a lot of scenes remain where only *either* the reflectance value *or* the curb information is available. In these cases the robustness of the road course extraction referring to different environments is additionally improved when using the fused representation of both features.

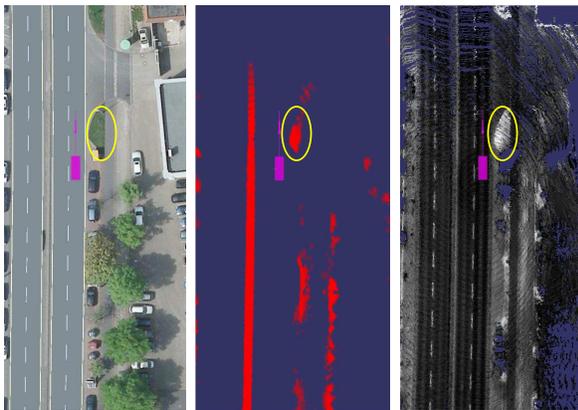


Fig. 9: Results of the curb feature extraction with redundant information: Left: Aerial view with digital map overlay; Middle: Extracted curb grid (height information); Right: Extracted reflectance grid (texture information). Aerial image: City of Braunschweig, Department of Geographic Information (no. 011/2010)

VII. CONCLUSION AND OUTLOOK

We have presented the extension of the vehicle environment modeling for automated road vehicles in urban scenarios by the usage of an explicit ground surface estimation. This estimation is done based on measurements of a 3D laser scanner, the Velodyne HDL-64 S2. Our approach provides locally consistent ground height estimations and is able to deal with slope changes of the ground surface inside the sensors' range. Moreover, it can deal with occlusions of the ground area and spatial false measurements. We have pointed out the enhancements based on the knowledge of the ground area, a pitch angle calculation and a curb feature extraction, which was shown to be an additional feature for road extraction algorithms. The estimated ground height is used to correct the height of object hypotheses and thus reduces the number of false positive and false negative object detections. Due to this, we were able to extend the sensors' effective field of view.

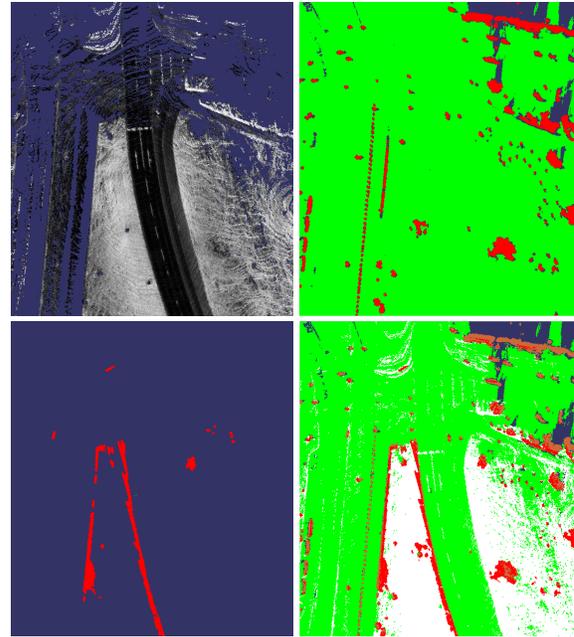


Fig. 10: Feature layers for the tristate road fusion. Upper left: reflectance data; upper right: elevated target occupancy; lower left: curb occupancy; lower right: resulting tristate fusion layer

At the current stage of research, the estimated height values stored in the grid structure are reinitialized with every new sensor scan to be able to react to high-frequent pitch and roll angle changes. Further research will address the storage of measured ground heights to fill temporarily seen cells not only based on its neighbor cells but also on past results from the ground plane estimation. This involves the compensation of pitch/roll influence in both timesteps and an efficient storage of the ground height during different measurement cycles.

The knowledge of seen ground surface areas allows the modeling of explicit free regions in the grid-based environment representation. Current implementations model the free regions in a Bayesian grid implicitly based on a free line-of-sight assumption up to the first elevated target. This assumption has some drawbacks which can lead to ambiguous interpretations: In case of sensor channels without valid measurements, it can either be interpreted as completely free field of view or as measurements of non-reflective objects. The explicit grid-based ground surface representation presented in this paper can be used to resolve these ambiguities in an easy way.

ACKNOWLEDGMENT

The authors wish to thank the Stadtpilot team for its support.

REFERENCES

- [1] J. M. Wille, F. Saust, and M. Maurer, "Stadtpilot: Driving autonomously on Braunschweig's inner ring road," *2010 IEEE Intelligent Vehicles Symposium (IV)*, pp. 506–511, Jun. 2010.
- [2] M. E. Bouzouraa, "Belegungskartenbasierte Umfeldwahrnehmung in Kombination mit objektbasierten Ansätzen für Fahrerassistenzsysteme," Ph.D. dissertation, Technische Universität München, 2011.

- [3] C. Guo, W. Sato, L. Han, S. Mita, and D. McAllester, "Graph-based 2D road representation of 3D point clouds for intelligent vehicles," *2011 IEEE Intelligent Vehicles Symposium (IV)*, pp. 715–721, Jun. 2011.
- [4] R. Matthaei, B. Lichte, and M. Maurer, "Robust grid-based road detection for ADAS and autonomous vehicles in urban environments," in *2013 16th International Conference on Information Fusion (FUSION)*, Istanbul, July 2013, pp. 938–944.
- [5] M. Kellner, M. Bouzouraa, and U. Hofmann, "Road curb detection based on different elevation mapping techniques," in *2014 IEEE Intelligent Vehicles Symposium Proceedings (IV)*, Dearborn, June 2014, pp. 1217–1224.
- [6] M. Montemerlo, J. Becker, S. Bhat, H. Dahlkamp, D. Dolgov, S. Ettinger, D. Haehnel, T. Hilden, G. Hoffmann, B. Huhnke, D. Johnston, S. Klumpp, D. Langer, A. Levandowski, J. Levinson, J. Marcil, D. Orenstein, J. Paefgen, I. Penny, A. Petrovskaya, M. Pflueger, G. Stanek, D. Stavens, A. Vogt, and S. Thrun, "Junior: The Stanford entry in the Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 9, pp. 569–597, Sep. 2008.
- [7] J. Bohren, T. Foote, J. Keller, A. Kushleyev, D. Lee, A. Stewart, P. Vernaza, J. Derenick, J. Spletzer, and B. Satterfield, "Little Ben: The Ben Franklin Racing Team's entry in the 2007 DARPA Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 9, pp. 598–614, Sep. 2008.
- [8] C. Urmsen, J. Anhalt, J. A. D. Bagnell, C. R. Baker, R. E. Bittner, J. M. Dolan, D. Duggins, D. Ferguson, T. Galatali, H. Geyer, M. Gittleman, S. Harbaugh, M. Hebert, T. Howard, A. Kelly, D. Kohanbash, M. Likhachev, N. Miller, K. Peterson, R. Rajkumar, P. Rybski, B. Salesky, S. Scherer, Y.-W. Seo, R. Simmons, S. Singh, J. M. Snider, A. T. Stentz, W. R. L. Whittaker, and J. Ziglar, "Tartan racing: A multi-modal approach to the darpa urban challenge," Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep., Apr. 2007.
- [9] J. Leonard, J. How, S. Teller, M. Berger, M. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman, O. Koch, Y. Kuwata, D. Moore, E. Olson, S. Peters, J. Teo, R. Truax, M. Walter, D. Barrett, A. Epstein, K. Maheloni, K. Moyer, T. Jones, R. Buckley, M. Antone, R. Galejs, S. Krishnamurthy, and J. Williams, "A Perception-Driven Autonomous Urban Vehicle," *Journal of Field Robotics*, vol. 1, no. 48, 2008.
- [10] S. Thrun and M. Montemerlo, "Stanley: The robot that won the DARPA Grand Challenge," *Journal of Field Robotics*, vol. 23, no. 9, pp. 661–692, 2006.
- [11] M. Himmelsbach and A. Müller, "LIDAR-based 3D object perception," *Proceedings of 1st International Workshop on Cognition for Technical Systems*, Oct. 2008.
- [12] J. Choi, S. Ulbrich, B. Lichte, and M. Maurer, "Multi-Target Tracking using a 3D-Lidar sensor for autonomous vehicles," *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, pp. 881–886, Oct. 2013.
- [13] M. Himmelsbach, F. V. Hundelshausen, and H.-J. Wuensche, "Fast segmentation of 3D point clouds for ground vehicles," *2010 IEEE Intelligent Vehicles Symposium (IV)*, pp. 560–565, Jun. 2010.
- [14] X. Hu, F. Rodriguez, and A. Geppert, "A multi-modal system for road detection and segmentation," in *2014 IEEE Intelligent Vehicles Symposium Proceedings (IV)*, June 2014, pp. 1365–1370.
- [15] C. Rodriguez-Garavito, A. Ponz, F. Garcia, D. Martin, A. de la Escalera, and J. Armingol, "Automatic laser and camera extrinsic calibration for data fusion using road plane," in *2014 17th International Conference on Information Fusion (FUSION)*, Salamanca, July 2014, pp. 1–6.
- [16] R. Matthaei, G. Bagschik, J. Rieken, and M. Maurer, "Stationary Urban Environment Modeling using Multi-Layer-Grids," in *17th International Conference on Information Fusion (FUSION 2014)*, Salamanca, 2014.
- [17] S. Perreault and P. Hébert, "Median Filtering in Constant Time," *IEEE Transactions on Image Processing*, vol. 16, no. 9, pp. 2389–2394, 2007.
- [18] J. Choi and M. Maurer, "Simultaneous Localization and Mapping based on the Local Volumetric Hybrid Map," in *2015 IEEE Intelligent Vehicles Symposium (IV)*, 2015.
- [19] D. Nuss, M. Thom, A. Danzer, and K. Dietmayer, "Fusion of laser and monocular camera data in object grid maps for vehicle environment perception," in *2014 17th International Conference on Information Fusion (FUSION)*, Salamanca, July 2014, pp. 1–8.
- [20] R. J. Adcock, "A problem in least squares," *The Analyst*, vol. 5, no. 2, pp. 53–54, 1878.
- [21] A. Seibert, M. Hahnel, A. Tewes, and R. Rojas, "Camera based detection and classification of soft shoulders, curbs and guardrails," in *2013 IEEE Intelligent Vehicles Symposium (IV)*, Gold Coast, June 2013, pp. 853–858.
- [22] J. S. Levinson, "Automatic laser calibration, mapping and localization for autonomous vehicles," Ph.D. dissertation, Stanford University, August 2011.
- [23] K. Dietmayer, J. Sparber, and D. Streller, "Model based object classification and object tracking in traffic scenes from range images," *2001 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1–6, 2001.
- [24] R. Matthaei, "Wahrnehmungsgestützte Lokalisierung in fahrstreifen-genauen Karten für Assistenzsysteme und automatisches Fahren in urbaner Umgebung," Ph.D. dissertation, Technische Universität Braunschweig, Braunschweig, (accepted to appear), 2015.